Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers

By Andrew Sweeting and Xuezhen Tao and Xinlu Yao*

We model repeated pricing by differentiated product firms when each firm has private information about its serially-correlated marginal cost. In a fully separating equilibrium of the dynamic game, signaling incentives can lead equilibrium prices to be significantly above those in a static, complete information game, even when the possible variation in the privately-observed state variables is very limited. We calibrate our model using data from the beer industry, and show that, without any change in conduct, our model can explain increases in price levels and changes in price dynamics and cost pass-through after the 2008 MillerCoors joint venture.
JEL: D43, D82, L13, L41, L66.
Keywords: oligopoly, dynamic games, dynamic pricing, asymmetric information, signaling, firm conduct, horizontal mergers, coordinated effects, merger retrospectives, beer industry.

Academic researchers and antitrust agency economists almost always use models that assume that firms set prices to maximize their profits in static and complete information (CI) environments. If an alternative is considered, it is typically tacit collusion with repeated CI stage games, often modeled by adding a conduct parameter term to static CI Nash first-order conditions. These formulations are tractable and, under fairly weak assumptions, the cost and conduct parameters are identified from price and market share data (Bresnahan (1982), Lau (1982), Nevo (1998), Berry and Haile (2014)).

However, the CI assumption that every firm knows everything that might affect its rivals' pricing choices appears inconsistent with how companies conceal the profitability of individual product lines, and with how antitrust agencies, even while using these models, presume that cost and margin data are confidential and competitively sensitive. A natural question is therefore whether predicted outcomes, and implications for antitrust enforcement, would change in a material way if the CI assumption is relaxed.

In this article, we consider non-collusive models where a small number of oligopolists repeatedly set prices, which are perfectly observable, and each firm has a privately observed, time-varying and positively serially-correlated state variable. We will focus on the case where this state variable is a firm's marginal cost, but, as illustrated in our working paper, Sweeting, Tao and Yao (2022) (STY-WP), results are similar when we allow another component of each firm's payoff function to be private information.

If higher marginal costs imply higher prices, and firms set higher prices and make higher profits

* Sweeting (corresponding author): University of Maryland, College Park, CEPR and NBER. Address: 3114 Tydings Hall, University of Maryland, College Park MD 20742, USA. Email: atsweet@umd.edu. Tao: Shanghai University of Finance and Economics. Address: 318 Wuchuan Rd, Wujiaochang, Yangpu District, Shanghai 200083, China. Email: tao.xuezhen@mail.shufe.edu.cn. Yao: c/o 3114 Tydings Hall, University of Maryland, College Park MD 20742, USA. Email: xlyao0420@gmail.com. We thank the Co-Editor (Robin Lee), three anonymous referees, Nate Miller, Matt Weinberg, Gloria Sheu, Dan Vincent, Bob Majure, Joe Harrington, Kyle Bagwell and a number of seminar participants for valuable comments, which have been especially useful in framing the paper. The empirical analysis uses the IRI Academic Dataset (Bronnenberg, Kruger and Mela (2008)), but the analysis is the work of the authors and not IRI. The predicted effects in Section III are from theoretical simulations and may not reflect the actual strategies of the firms considered in this paper. This research was supported by US National Science Foundation Grant SES-1658670 (Sweeting) and National Natural Science Foundation of China Grant 72003120 (Tao). Xinlu Yao's work on this project was completed when she was a graduate student at the University of Maryland. Qian Wang provided research assistance for the demand estimation. All errors are our own.

when their rivals set higher prices, firms will have an incentive to raise their prices in order to signal that they have higher current costs. We consider fully separating equilibria where each firm can infer its rivals' current marginal costs from the prices that they set. The simplicity of equilibrium beliefs leads to computational tractability.[1]

In a two-period model with linear demand, Mailath (1989) shows that signaling will raise equilibrium prices relative to the CI model. The novelty in our article comes from computing the size and the policy implications of these effects.[2] In games with longer time horizons, signaling can raise prices significantly (i.e., by more than a few percentage points) and, because both merging and non-merging firms can have stronger incentives to signal after a horizontal merger (that does not lead to monopoly), our model can predict substantially larger post-merger price increases in both symmetric and asymmetric industries. Importantly, the effects we identify can be large even when the degree of private information is small, in the sense that the supports of marginal costs are narrow. Therefore, one only needs to slightly relax the CI assumption to get substantively different predictions.

In asymmetric industries, we find examples where signaling leads non-merging firms to raise their prices as much or more than the merged firm. This partly motivates our empirical application where we calibrate a stylized version of our model using data from the U.S. beer market prior to the 2008 MillerCoors (MC) joint venture (JV). Miller and Weinberg (2017*b*) (MW) show that MC and its larger domestic rival Anheuser-Busch (AB), raised their prices by similar amounts after the JV. MW use the fact that this cannot be rationalized by a static CI Nash model to identify a collusive conduct parameter.

When we calibrate our model using pre-JV data and make plausible assumptions about likely synergies, we find that our model predicts increases in price levels that are similar to those observed in the data, and that our model also predicts some qualitative changes in price dynamics that are not predicted by a simple conduct parameter model. Under some additional assumptions about which aspect of marginal cost is private information, our model can also explain some observed changes in the pass-through of transportation costs. While folk theorems suggest that there is likely to be a richer collusive model that could also explain the changes in price dynamics and pass-through, we interpret our results as showing that one does not need to assume collusion to explain why post-merger prices increase more than CI models predict (Peters (2009), Garmon (2017), Ashenfelter, Hosken and Weinberg (2014)).[3]

After a review of the related theoretical literature, Section I presents our model and equilibrium concept. Section II presents examples and illustrates the implications for merger analysis. Section III provides our empirical application. Section IV concludes. The online appendices detail our computational algorithms, some additional examples (with more explored in STY-WP) and our data. Code to replicate our results is available online (Sweeting, Tao and Yao (2023)).

Related Literature. — Mailath (1988) identifies conditions under which a separating equilibrium will exist in an abstract two-period game with continuous types, and shows that the conditions on payoffs required for the uniqueness of each player's separating best response functions are similar to the ones required in a model where one agent signals. Mailath (1989) applies these results to a

[1] Fershtman and Pakes (2012) and Asker et al. (2020) develop dynamic asymmetric information games with discrete actions, and reduce the computational burden by assuming that players maintain beliefs about expected payoffs from different actions rather than rivals' types.

[2] Shapiro (1986) and Vives (2011) show how asymmetric information about costs directionally affects prices and welfare in one-shot models, where there are no signaling incentives. For the parameters we consider, the differences between static CI prices and static Bayesian Nash equilibrium prices are tiny.

[3] In a three period quantity-setting model, Mester (1992) shows that marginal cost signaling will increase output, relative to a CI Cournot model, as firms want their rivals to produce less. Mester's model is motivated by trying to explain outcomes in banking where more concentration appeared to lead to more competitive behavior, a result consistent with the banking merger retrospectives cited in Ashenfelter, Hosken and Weinberg (2014).

two-period duopoly pricing game with linear demand, and shows that there is a unique equilibrium where signaling increases first-period prices. We rely on Mailath's results to characterize best response signaling pricing functions, and use computation to extend this model to more periods, more firms and a richer demand structure. However, we cannot claim equilibrium uniqueness.

A feature of our model is that marginal costs can evolve over time, so there is always an incentive to signal except in the last period. Kaya (2009) and Toxvaerd (2017) consider repeated games where one firm has a privately known state that is fixed, and it signals until its reputation is established. Harrington (2021) and Sweeting, Leccese and Tao (2023) consider models where a merged firm has private information on the size of a realized marginal cost synergy, and firms may pool on the prices that would form a CI Nash equilibrium only for the smallest possible realization of the synergy.

Our article is also related to Sweeting, Roberts and Gedge (2020) (SRG). SRG develop dynamic versions of the Milgrom and Roberts (1982) limit pricing model where an incumbent monopolist's cost evolves over time so that there is persistent price signaling, and they show that the model explains why incumbent airlines lowered prices by 15% on some routes when Southwest threatened entry. The differences are that the current paper considers more common oligopoly market structures and a more broadly relevant policy application. The effects are also large for different reasons. In SRG, an incumbent may lower its price significantly because entry by an efficient rival will permanently and dramatically lower its profits. In the current paper, any signal only affects outcomes one-period ahead, and it is feedbacks between the signaling strategies of different firms that generate large equilibrium effects.

I. Model

A. *Specification.*

A fixed set of N risk-neutral firms simultaneously set prices in each period of a game with $t = 1, ..., T$ periods, where $T \leq \infty$. Unless otherwise stated, the discount factor is $\beta = 0.99$. If a firm sells multiple products, we will assume that those products are symmetric and must be sold at a common price. There may be commonly known differences in demand and marginal costs across firms, but exactly one dimension of a firm's type is private information. All firms observe current and past prices.

We will consider two different formulations of types. The first formulation ("continuous types") assumes firm i's type can take any value on a known compact interval $[\underline{\theta_i}, \overline{\theta_i}]$, whereas the second formulation ("discrete types") assumes i's type is either $\underline{\theta_i}$ or $\overline{\theta_i}$. The second formulation is useful when we need to lower the computational burden.

In both cases, types evolve exogenously and independently according to a first-order Markov process, $\psi_i : \theta_{i,t-1} \rightarrow \theta_{i,t}$. In this article, we will assume that types reflect marginal costs, in which case the transition assumption, while strong, is consistent with the assumptions in Olley and Pakes (1996) and the subsequent structural production function literature.

Within-Period Timing. — In each period t, timing is as follows. Firms enter t with their $t-1$ types, which then evolve according to ψ_i. Each firm observes its own new type, but not the previous or new type of rivals.[4] Each firm simultaneously chooses (and commits to) its price, $p_{i,t}$, with no menu costs. Demand is static and time-invariant. i's one period profit is $\pi_i(p_{i,t}, p_{-i,t}, \theta_{i,t})$ and we assume that $\frac{\partial \pi_i}{\partial p_{-i,t}} > 0$ for all $-i$.

[4] Our fully separating equilibria would be unchanged if $t-2$ types are revealed in t.

TRANSITION ASSUMPTIONS. —

Assumption 1 ***Type Transitions for the Continuous Type Model.*** *The conditional pdf* $\psi_i(\theta_{i,t}|\theta_{i,t-1})$

(i) has full support, so that the type can transition from any value on the support to any other value in a single period.

(ii) is continuous and differentiable (with appropriate one-sided derivatives at the boundaries).

(iii) for any $\theta_{i,t-1}$ *there is some* θ' *such that* $\frac{\partial\psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial\theta_{i,t-1}}|_{\theta_{i,t}=\theta'} = 0$ *and* $\frac{\partial\psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial\theta_{i,t-1}} < 0$ *for all* $\theta_{i,t} < \theta'$ *and* $\frac{\partial\psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial\theta_{i,t-1}} > 0$ *for all* $\theta_{i,t} > \theta'$*.*

Assumption 1(iii) implies that each firm's marginal cost is positively serially correlated in the sense that a higher cost in period t implies a higher cost in $t+1$ is more likely. In the discrete type model we will assume there is a probability, $\rho > 0.5$, that each firm's cost remains the same in the next period. Assumption 1(i) is important because it implies that a firm will keep on having incentives to signal its current marginal cost even in a long game.

We will consider fully separating equilibria where each firm has correct beliefs, in equilibrium, about each rival's previous period type. In order to avoid the inconvenience of calculating different strategies for the first period of a game, we make the following assumption so that beliefs in $t = 1$ have the same structure.

Assumption 2 ***Initial Period Beliefs.*** *Firms know what their rivals' types were in a fictitious prior period,* $t = 0$*.*

B. Fully Separating Equilibrium in a Finite Horizon and Continuous Type Game.

We now describe the equilibrium for a continuous-type game with two ex-ante symmetric single-product duopolists. We start in the last period (T), assuming that play in $T-1$ was separating so that each player has a point belief about the $T-1$ costs of other players.

FINAL PERIOD (T). — In the final period, firms use Bayesian Nash Equilibrium (BNE) strategies that maximize their expected payoffs given their own types, their beliefs about their rival's type and their rival's strategy. If firm j believes that firm i's period $T-1$ type was $\widehat{\theta^j_{i,T-1}}$ and j's period T pricing function is $P_{j,T}(\theta_{j,T}, \theta_{j,T-1}, \widehat{\theta^j_{i,T-1}})$[5], then a type $\theta_{i,T}$ i will set a price

(1)
$$p^*_{i,T}(\theta_{i,T}, \theta_{j,T-1}, \widehat{\theta^j_{i,T-1}}) = \arg\max_{p_{i,T}} \int_{\underline{\theta_j}}^{\overline{\theta_j}} \pi(p_{i,T}, P_{j,T}(\theta_{j,T}, \theta_{j,T-1}, \widehat{\theta^j_{i,T-1}}), \theta_{i,T})\psi(\theta_{j,T}|\theta_{j,T-1})d\theta_{j,T}.$$

EARLIER PERIODS ($1, .., T-1$). — In earlier periods, i may choose not to set a static best response price because its price can affect future play, via rivals' beliefs. The equilibrium concept that we use is symmetric Markov Perfect Bayesian Equilibrium (MPBE) (Toxvaerd (2008), Roddie (2012)). An MPBE specifies each firm's beliefs about costs given observed prices and period-specific pricing

[5]This notation reflects the fact that we are assuming that player j used an equilibrium strategy in $T-1$ that revealed its type ($\theta_{j,T-1}$), but we are allowing for the possibility that firm i may have deviated so that j's beliefs about i's previous type are incorrect.

strategies for each firm i as a function of its current type and beliefs, where the strategy maximizes i's discounted future payoff given rivals' strategies. History can matter because it affects beliefs.

In a fully separating MPBE, each firm's pricing function allows rivals to infer the firm's current cost from its chosen price, and equilibrium beliefs have a simple form. For completeness, we also need to define beliefs that a firm will have if the rival sets a price that is outside the range of the equilibrium pricing function. As price functions are continuous, we will assume that when a firm sets a price below (above) the lowest (highest) price in the range of the pricing function, it will be inferred to have the lowest (highest) possible cost type.

Characterization of Separating Pricing Functions in Period $t < T$. — We characterize fully separating pricing functions by defining a firm i's period-specific "signaling payoff function", $\Pi^{i,t}(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t})$, following Mailath (1989). $\Pi^{i,t}$ is the present discounted value (PDV) of firm i's expected current and future payoffs when its current type is $\theta_{i,t}$, it sets price $p_{i,t}$ and j believes, *at the end of period* t, that i has type $\widehat{\theta^j_{i,t}}$. $\Pi^{i,t}$ is assumed to be continuous and at least twice differentiable in its arguments. It is implicitly conditional on (i) j's period t pricing strategy, which will depend on j's beliefs about $t-1$ types, and (ii) both players' strategies in future periods. As j's end-of-period t belief about i's type enters as a separate argument, $p_{i,t}$ only affects $\Pi^{i,t}$ through period t profits.

Following Mailath, the separating best response function of firm i, which is also implicitly conditioned on j's current pricing strategy and beliefs about previous types, can, under the conditions to be listed in a moment, be uniquely characterized as follows: i's pricing function will be the solution to a differential equation where

$$\frac{\partial p^*_{i,t}(\theta_{i,t})}{\partial \theta_{i,t}} = -\frac{\Pi^{i,t}_2\left(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t}\right)}{\Pi^{i,t}_3\left(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t}\right)} > 0, \tag{2}$$

and a boundary condition. The subscript n in $\Pi^{i,t}_n$ denotes the partial derivative of $\Pi^{i,t}$ with respect to the n^{th} argument. Assuming that lower types want to set lower prices (e.g., a type corresponds to the firm's marginal cost), the boundary condition will be that $p^*_{i,t}(\underline{\theta_i})$ is the solution to

$$\Pi^{i,t}_3\left(\underline{\theta_i}, \widehat{\theta^j_{i,t}}, p_{i,t}\right) = 0, \tag{3}$$

i.e., the lowest type's price maximizes its static expected profits given j's pricing policy. The numerator in (2) is i's marginal future benefit from raising j's belief about $\theta_{i,t}$, and the denominator is the marginal effect of a price increase on i's current profit. For prices above a static best response price, the denominator will be negative, and the pricing function will slope upwards in the firm's type.

$\Pi^{i,t}$ must satisfy the following conditions for this characterization to hold.

Condition 1 ***Shape of*** $\Pi^{i,t}$ ***with respect to*** $p_{i,t}$. *For any* $(\theta_{i,t}, \widehat{\theta^j_{i,t}})$, $\Pi^{i,t}\left(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t}\right)$ *has a unique optimum in* $p_{i,t}$, *and, for all* $\theta_{i,t}$, *for any* $p_{i,t}$ *where* $\Pi^{i,t}_{33}\left(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t}\right) > 0$, *there is some* $k > 0$ *such that* $\left|\Pi^{i,t}_3\left(\theta_{i,t}, \widehat{\theta^j_{i,t}}, p_{i,t}\right)\right| > k$.

Condition 2 ***Type Monotonicity.*** $\Pi_{13}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right) \neq 0$ *for all* $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t})$*.*

Condition 3 ***Belief Monotonicity.*** $\Pi_{2}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)$ *is either* > 0 *for all* $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}})$ *or* < 0 *for all* $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}})$*.*

Condition 4 ***Single-Crossing.*** $\frac{\Pi_{3}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)}{\Pi_{2}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)}$ *is a monotone function of* $\theta_{i,t}$ *for all* $\widehat{\theta_{i,t}^{j}}$ *and for* $(\theta_{i,t}, p_{i,t})$ *in the graph of* $p_{i,t}^{*}(\theta_{i,t}, \theta_{j,t-1})$*.*

Assuming types correspond to marginal costs, the first condition requires that a firm's current expected profit is always quasi-concave in its own price, whatever prices rivals set, for all possible costs. This is true for common forms of demand, such as the nested logit model. Type monotonicity requires only that, when a firm increases its price, its (current) lost profit is smaller when its marginal cost is higher. Belief monotonicity requires that a firm's expected future profits should increase when rivals believe its current cost is higher (holding its actual cost fixed) and single-crossing requires that a higher marginal cost firm is more willing to raise its price in order to increase rivals' beliefs about its cost. Online Appendix B.4 uses a discrete type example to illustrate why belief monotonicity and single-crossing can fail when prices rise too much from static CI Nash levels.

These conditions are sufficient for a fully separating *best response function* to exist and be unique, but they do not imply the existence or uniqueness of a fully separating *equilibrium* when multiple firms are signaling. This is the case even when the linear demand duopoly model of Mailath (1989) is extended to more than two periods.[6]

We use computation to find an equilibrium (see online Appendix A for details), assuming that pricing functions have the form described above and verifying that the required conditions hold at our solutions. This approach allows us to consider longer games, more firms and non-linear demand systems. In games with continuous types, we have found that using different starting points and updating rules for strategies leads to equilibrium pricing functions that are almost identical.

C. Fully Separating Equilibrium in a Finite Horizon and Discrete Type Game.

In a discrete type game, even separating best responses may not be unique. As described in online Appendix B, our solution method therefore applies a refinement by always solving for the best response strategies that achieve separation at the lowest cost to the signaling firm given our current guess of rivals' strategies. This is consistent with the type of "intuitive criterion" (Cho and Kreps (1987)) refinement that has been widely used in one-sided signaling models with two types. However, we have identified examples of multiplicity in discrete type games with infinite horizons. In the text, we therefore report results when we solve finite horizon games and iterate backwards until strategies converge, although, even in this case, there is no guarantee of uniqueness.

II. Examples

We now use examples to illustrate the effects of simultaneous signaling on prices and on the effects of mergers. Appendix B of STY-WP provides further examples where marginal costs are fixed and known, and some other element of each firm's profit function is private information.

[6]Appendix C of STY-WP shows that existence and uniqueness require that the effects of signaling on prices are sufficiently small to guarantee single-crossing.

A. A Continuous-Type Symmetric Duopoly Example.

SPECIFICATION. — Two ex-ante symmetric single-product firms play a finite horizon game with $T \geq 25$ periods. Demand is nested logit, with both products in one nest and the outside good in its own nest. Consumer c's indirect utility from buying product i is $u_{i,c} = 5 - 0.1p_i + \sigma\nu_c + (1-\sigma)\varepsilon_{i,c}$. where p_i is the dollar price, $\varepsilon_{i,c}$ is a draw from a Type I extreme value distribution and ν_c is an appropriately distributed draw for c's nest preferences (Cardell (1997)), given a nesting parameter of $\sigma = 0.25$. For the outside good, $u_{0,c} = \varepsilon_{0,c}$.

Firm marginal costs are private information. For each firm, $c_{i,t}$ lies in the interval $[\underline{c}, \overline{c}] =$ [\$8, \$8.05] and evolves according to an independent and exogenous truncated AR(1) process

$$c_{i,t} = \rho c_{i,t-1} + (1-\rho)\frac{\overline{c}+\underline{c}}{2} + \eta_{i,t}, \tag{4}$$

where $\eta_{i,t} \sim TRN(0, \sigma_c^2, \underline{c} - \rho c_{i,t-1} - (1-\rho)\frac{\overline{c}+\underline{c}}{2}, \overline{c} - \rho c_{i,t-1} - (1-\rho)\frac{\overline{c}+\underline{c}}{2})$.[7] $\rho = 0.8$ and $\sigma_c = \$0.025$.

The demand parameters and the level of marginal costs imply that CI Nash margins are high, with limited demand diversion to the outside good. As illustrated in online Appendix B.3, limited diversion plays an important role in supporting separating equilibria with significant price effects.

On the other hand, the cost transition parameters imply that marginal costs can only vary by a small amount and that, within the allowed range, they are quite likely to vary from being relatively high to relatively low (or vice-versa) from one period to the next. One summary measure of persistence is $\Pr(c_{i,t+1} \geq \frac{\overline{c}+\underline{c}}{2} | c_{i,t} = \overline{c})$, which equals 0.68 for our parameters. This implies that j cannot gain much information on the likely value of $c_{i,t+1}$ when it infers $c_{i,t}$, so that i's incentives to distort its price to signal in period t should be limited.

EQUILIBRIUM STRATEGIES AND OUTCOMES. — In period T, firms use one-shot BNE pricing strategies. Figure 1(a) shows four pricing functions for firm 2, for different values of firm 1's period $T-1$ marginal cost ($c_{1,T-1}$), assuming that both firms know/believe that $c_{2,T-1} = \$8$. Firm 2's price increases with $c_{1,T-1}$ because firm 1's expected period T price rises with $c_{1,T-1}$. However, the variation in $c_{1,T-1}$ can change $p_{2,T}$ by no more than one cent. Average prices and welfare measures are almost identical to static CI Nash outcomes, implying that the effects of incomplete information with no dynamics (Shapiro (1986), Vives (2011)) are negligible given our parameters.[8]

In period $T-1$, a firm's price may affect its rival's period T price, creating a signaling incentive. Assuming both firms' period $T-2$ costs were \$8, Figure 1(b) shows firm 1's signaling pricing function (found by solving the differential equation (2) given the boundary condition (3)) if (hypothetically) it expects firm 2 (in period $T-1$) to use its static period T strategy. The period T pricing strategies are shown for comparison. The pricing functions intersect for $c_{1,T-1} = \$8$, but signaling may lead firm 1 to raise its price by as much as 20 cents for higher $c_{1,T-1}$ values.

This strategy is profitable because firm 1's period $T-1$ profit is quite flat in its own price. Figure 2 shows that if $c_{1,T-1} = \$8.025$, setting the signaling price of \$22.76, rather than the statically optimal price, only lowers $T-1$ profits by \$0.00070 (per potential consumer). This is smaller than the expected period T gain (\$0.00079) from being viewed as a firm with $c_{1,T-1} = \$8.025$ rather than $c_{1,T-1} = \$8.0002$ (firm 2's inference if firm 1 set its statically optimal price).

Of course, firm 2 also has an incentive to signal. Figure 1(c) shows firm 2's best signaling response when (hypothetically) firm 1 uses the strategy in Figure 1(b) (once again shown for comparison). As firm 1's expected price is rising, firm 2's static best response pricing function shifts upwards.

[7] TRN denotes a truncated normal distribution, whose arguments are the mean and the variance of the untruncated distribution, and the lower and upper truncation points.

[8] For example, expected consumer surplus differs by less than \$0.0001 per potential consumer.

Figure 1. : Period T and $T-1$ Pricing Strategies in the Finite Horizon, Continuous Type Signaling Game.

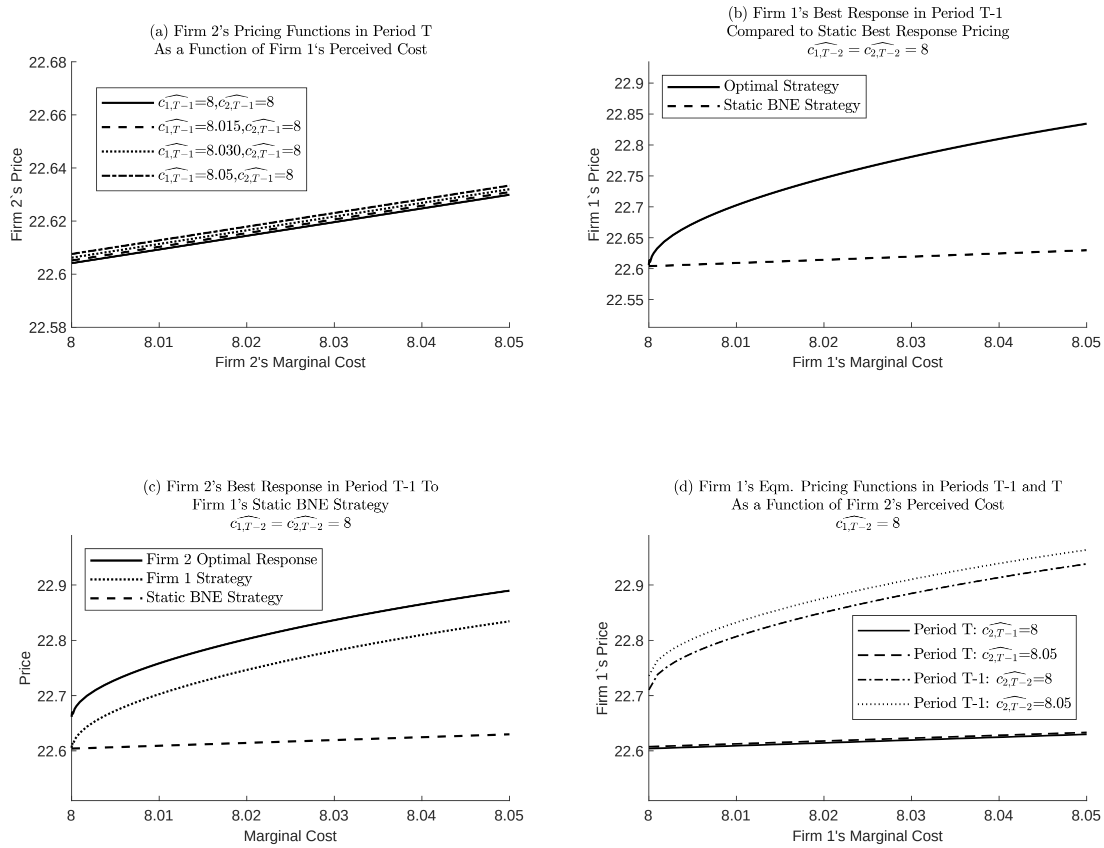

Of course, this will, in turn, affect firm 1's pricing function. Figure 1(d) shows the equilibrium period $T-1$ pricing functions where firms are best responding to each other. The price functions are steeper and more spread out than those for period T, implying that prices will have higher variance as well as a higher average level.[9]

While the qualitative effects of signaling on $T-1$ prices are similar to those considered by Mailath (1989), we are interested in what happens to prices in longer games. As the $T-1$ price functions are more spread out, signaling incentives will tend to be greater in period $T-2$ than period $T-1$. As a result, Figure 3 shows that $T-2$ pricing functions are higher and have steeper slopes than those in $T-1$. They are also more spread out, leading to even stronger incentives in $T-3$. As we continue backwards, the pricing functions continue to rise until around $T-15$, after which they change only slightly. In $T-14$ average prices are \$24.76 (or \$2 above static CI Nash levels) and the

[9]The mean (standard deviation) of a firm's equilibrium price in period $T-1$ is 22.88 (0.065), compared with 22.62 (0.003) in period T.

Figure 2. : Expected $T-1$ Period Profit Function: $c_{1,T-1} = \$8.025$ and $c_{1,T-2} = c_{2,T-2} = \$8$.

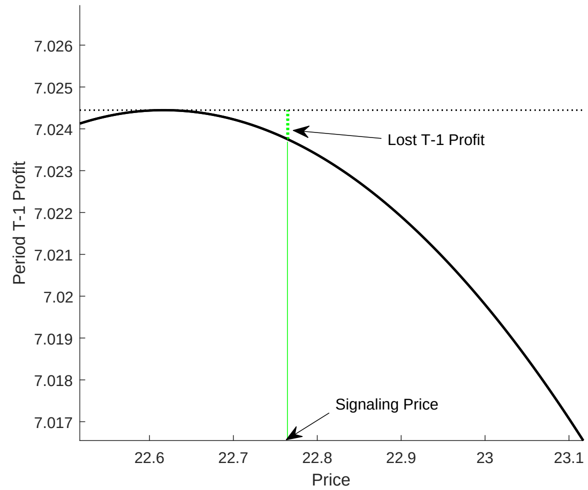

Note: the profit function is drawn "per potential consumer" for a firm assumed to have a marginal cost of \$8.025, and with a rival using the static BNE pricing strategy when both firms' previous period marginal costs were \$8.

standard deviation of each firm's price is \$0.47. Prices are, however, much lower than joint-profit maximizing CI prices (\$45.20).

We can also solve for stationary pricing functions in an infinite horizon game. In this case, we calculate firms' best responses when they expect that the current guess of the pricing functions will be used in all future periods. The computed pricing strategies are visually indistinguishable from the strategies in the early periods of the $T = 25$ finite horizon game, although our calculations imply that average prices may be a cent higher.[10]

B. *Mergers with Continuous Types.*

We can extend the example, maintaining the same demand and cost parameters, to consider the effect of mergers. Specifically, we consider a 4-to-3 merger and a 3-to-2 merger, where, after the merger, one of the firms owns two products that have the same marginal cost realization each period and for which the seller must set exactly the same price. In this way, the one unobserved state variable-one signal per firm structure is maintained.[11]

Table 1 presents the results, based on infinite-horizon signaling prices. When there are three single-product firms, average prices are \$19.89. This contrasts with expected static BNE prices of \$19.20. Absent cost efficiencies, a 3-to-2 merger raises the merged firm's average price by 38% and the average price of the non-merging firms by 20%. With static BNE (or static CI Nash) pricing, the price increases would be 28% and 10% respectively. The predicted price increases after a 4-to-3 merger are also substantial (in this case, merging and non-merging firm price increases would be 16% and 3% with static BNE pricing).

Figure 4 illustrates how pricing strategies change for a merging party and a non-merging party in the static BNE and infinite horizon dynamic signaling games after a 3-to-2 merger. In the BNE

[10]We do not interpret this difference as significant, as we have found that the number of gridpoints, the choice of methods for integrating and solving differential equations and convergence criteria can change average prices by up to 5 cents in long games. We have found that finite horizon strategies converge so that they are close to the infinite horizon ones in all of our continuous type examples, unless Mailath's conditions are very close to being violated, in which case the finite horizon strategies often oscillate up and down across periods.

[11]In these examples, we will assume that, if a synergy is realized, the magnitude of the synergy is known to all players. Sweeting, Leccese and Tao (2023) considers a game where the magnitude of the realized synergy is private information to the merging firms.

Figure 3. : Equilibrium Pricing Functions for Firm 1 in the Infinite Horizon Game and Various Periods of the Finite Horizon Game.

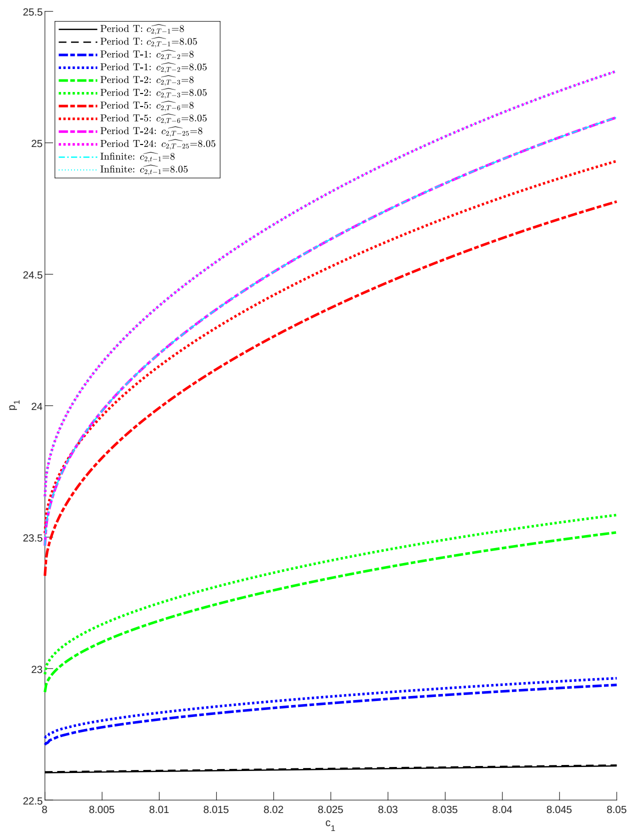

Note: all functions drawn assuming that firm 1's perceived marginal cost in the previous period was \$8.

case, the merged firm internalizes diversion between its products, which causes its optimal price to increase. The non-merging firm's best response price function shifts upwards as a result. In the signaling case, the pricing functions of both the merged firm and the non-merging firm also become steeper and more spread out.[12] This reflects how the elimination of the third firm leads to both surviving firms having a rival whose future price will be more responsive to the price that it sets, leading signaling incentives to increase.

Typically, an antitrust agency would challenge a 4-to-3 or 3-to-2 merger unless they expect the merged firm to realize significant marginal cost efficiencies. Agencies usually calculate the Compensating Marginal Cost Reduction implied by a static CI Nash model (CI CMCR), to assess whether likely efficiencies are "large enough". Table 1 reports the CI CMCRs an analyst would calculate if she knows demand and calculates pre-merger costs assuming that average pre-merger prices reflect static CI Nash pricing.[13]

[12]In contrast, the BNE pricing functions become slightly flatter after the merger as the markups increase.

[13]For example, when there are three firms before the merger, the analyst would calculate that each pre-merger firm has a marginal cost of \$8.73, whereas the true average cost is \$8.025. For a 3-to-2 merger, a CI CMCR synergy would give the merged

Table 1—: Post-Merger Prices and Required Synergies in an Infinite Horizon Continuous-Type Model. Firms are ex-ante symmetric before the merger. After the merger, the merged firm sells two products that have identical cost realizations and the same price.

	4-to-3 Merger	3-to-2 Merger
Pre-Merger Average Prices	\$18.24	\$19.89
Post-Merger Average Price of Merged Firm if No Marginal Cost Synergy	\$21.61 (+18.4%)	\$27.46 (+38.1%)
Post-Merger Average Price of Non-Merging Firm if No Marginal Cost Synergy	\$19.21 (+5.3%)	\$23.85 (+19.9%)
CI CMCR	\$4.87	\$10.65
Merged's Firm Post-Merger Average Price with CI CMCR Synergy in Signaling Equilibrium	\$18.75 (+2.8%)	\$23.14 (+16.4%)
Marginal Cost Reduction Required to Keep Merged Firm's Average Price from Rising in Signaling Equilibrium	\$5.81	\$21.10

Note: parameterization described in the text. Note that the CI CMCR is the marginal cost reduction that an analyst would compute using the true demand system, observed (signaling) pre-merger signaling prices and a CI Nash assumption.

The CI CMCRs for these mergers are large, but even if they are realized, post-merger prices would increase if firms use signaling strategies.[14] For example, the prices of the merged firm and the non-merging firm would rise by 16% and 13% after a 3-to-2 merger. The final row of Table 1 reports the marginal cost reductions required to keep the merged firm's average price at its pre-merger level in our model. This required efficiency is 19% larger than the CI CMCR in the 4-to-3 case, and more than 100% larger in the 3-to-2 case.[15]

C. Alternative Specifications.

Changing Cost Parameters for Duopoly. — We have suggested that the values of σ_c and $\overline{c} - \underline{c}$ imply that signaling incentives are relatively weak in our baseline specification, as signaling a high cost in period t cannot have too much effect on the cost that the rival expects in $t+1$.[16] We might therefore expect even larger price effects if we lower σ_c or raise $\overline{c} - \underline{c}$.[17]

The second to sixth columns of Table 2 show that this intuition is correct, but that, when prices increase too much, the conditions required to characterize best response functions fail.[18]

firm a negative marginal cost. The analyst assumes that the merged firm cannot freely dispose of the good, so that it would not choose to produce an infinite amount.

[14]We assume that both $\underline{c}$ and $\overline{c}$ for the merged firm fall by the CI CMCRs with the other cost parameters unaffected.

[15]These efficiencies are still not large enough to prevent consumers from being harmed as the non-merging firm's price would still increase (for example, by an average of \$1.61 in the 3-to-2 case).

[16]For example, the expected $c_{i,t+1}$s if $c_{i,t}$ takes on its minimum and maximum values are 8.0194 and 8.0306 respectively.

[17]The logic for why each of these changes would raise prices is slightly different. A reduction in σ_c makes costs more persistent, so that signaling a higher current cost will have a greater effect on rivals' expectations of next period's costs. This will tend to raise the numerator in the differential equation, causing the pricing functions to become steeper. An increase in $\overline{c}$, holding $\underline{c}$ fixed, will tend to raise each firm's average signaling price by more than their average static BNE price if the signaling pricing functions are steeper than the static BNE functions. In both cases, these changes will cause rivals' best response prices to rise, and their price functions to be translated upwards.

[18]Of course, equilibria that involve different types of strategies may exist, but we do not know how to find them.

Figure 4. : Static Bayesian Nash and Infinite Horizon Dynamic Signaling Equilibrium Pricing Functions Before and After a 3-to-2 Merger.

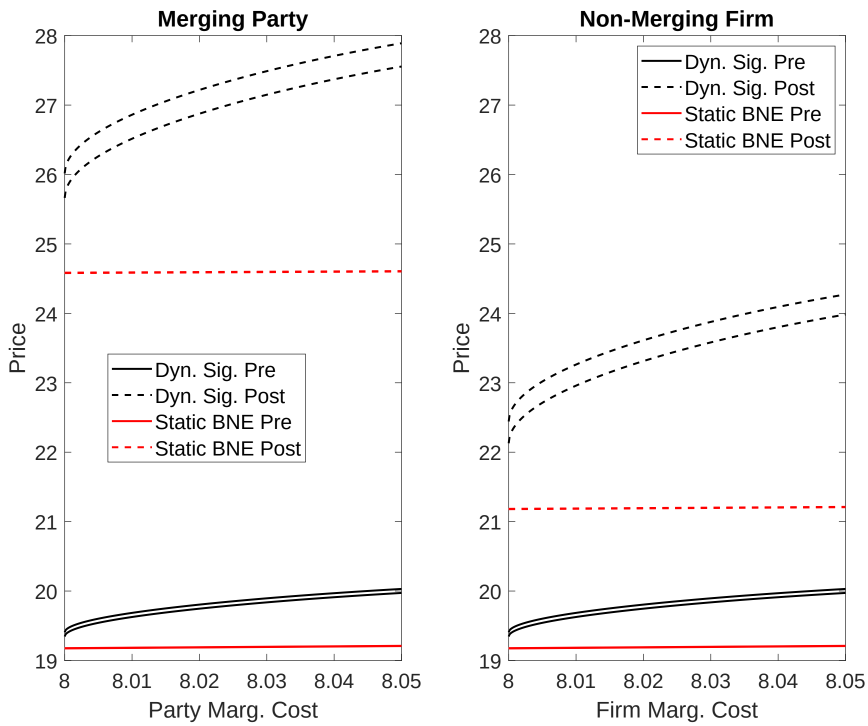

Note: parameters discussed in the text. Pricing functions shown are those when all products have marginal costs of 8 in the previous period, and those where all products have marginal costs of 8.05 in the previous period, although these different pricing functions cannot be visually distinguished for the static BNE cases.

The final column increases both $\bar{c}$ and σ_c so that we maintain $\Pr(c_{i,t+1} \geq \frac{\bar{c}+\underline{c}}{2}|c_{i,t} = \bar{c}) \approx 0.68$. In this case, a fully separating equilibrium exists with a similar price increase to our baseline specification. This example is relevant for thinking about the calibrated parameters in our empirical application.

These examples assume that the duopolists are symmetric. In Appendix C.6, we will also discuss some assumptions which would lead to the supports of marginal costs of duopolists being different. While we have not explored these types of examples systematically, in the one or two examples we have looked at we have found that increases in the average prices of both firms are more sensitive to the narrower of the cost supports.[19]

Additional Firms and Alternative Discount Factors. — We use the two-type model to consider a wider range of market structures and alternative discount factors.

First, we vary the number of symmetric single-product firms from 2 to 7, and consider discount factors of 0.8, 0.95 and 0.99. Each firm's marginal cost state is either 8 or 8.05. The Markov probability that a firm's cost remains the same in the next period is $\rho = 0.65$. The demand

[19]For example, if we increase the $\bar{c}$ of one firm from 8.05 to 8.075, and leave the other fixed, the average prices of both firms in a long or infinite horizon game are similar to those in the first column, whereas when we increase the $\bar{c}$ of both firms, average prices increase significantly (second column).

Table 2—: Equilibrium Expected Prices in a Finite Horizon Game with Alternative Cost Specifications.

	Baseline	Expand Range			Reduce Std. Deviation		Expand Range & Increase Std. Dev.
$[\underline{c}, \overline{c}]$ ($)	[8,8.05]	[8,8.075]	[8,8.15]	[8,8.3]	[8,8.05]	[8,8.05]	[8,8.50]
σ_c ($)	0.025	0.025	0.025	0.025	0.02	0.01	0.25
Static CI Nash Pricing							
Every Period	$22.62	$22.63	$22.67	$22.74	$22.62	$22.62	$22.84
Dynamic Signaling Game							
T-24	$24.70	$26.43	-	-	$25.63	-	$24.87
T-10	$24.69	$26.50	-	-	$25.62	fails	$24.87
T-9	$24.68	$26.50	fails	-	$25.60	$27.94	$24.86
T-8	$24.66	$26.47	$28.21	-	$25.58	$28.44	$24.86
T-7	$24.62	$26.41	$28.98	fails	$25.52	$28.63	$24.84
T-6	$24.55	$26.28	$29.18	$29.85	$25.41	$28.10	$24.82
T-3	$23.85	$24.88	$26.89	$28.33	$24.31	$25.93	$24.47
T-2	$23.37	$23.95	$25.11	$26.21	$23.60	$24.33	$24.12
T-1	$22.88	$23.05	$23.43	$23.94	$22.93	$23.05	$23.56
T	$22.62	$22.63	$22.67	$22.74	$22.62	$22.62	$22.84
∞-Horizon	$24.71	$26.45	fails	fails	$25.64	fails	$24.88

Note: values in all but the last line are based on the duopoly, continuous type, finite horizon model with demand parameters described in the text (cost parameters indicated in the table). The last line reports results for the stationary strategies in the infinite horizon model with the same parameters. "Fails" indicates that the belief monotonicity or single-crossing conditions fail so that we cannot calculate signaling best response pricing functions.

parameters are the same as in our continuous type example. As we have identified multiple equilibria in the infinite horizon discrete type game, we solve a finite horizon game backwards for at least 30 periods, adding additional periods, up to 150, until the pricing strategies change by less than 5e-4 across periods.

Figure 5(a) shows how much prices increase in the dynamic signaling equilibrium, relative to static CI Nash prices.[20] There is trivially no effect in the monopoly case as there is no rival to signal to. The effect of signaling declines in the number of firms and increases in the discount factor. When $N = 4$, price increases are more than 1% even with a discount factor consistent with annual pricing.

Figure 5(b) uses the same computations to show the predicted price increases after a two firm merger that results in one of the products being eliminated and no synergies (so that the firms remain symmetric). The static CI Nash model naturally predicts a larger effect for a merger to monopoly, but for the other mergers, dynamic signaling results in larger price effects, with non-trivial differences to the static CI Nash model after 3-to-2, 4-to-3 and 5-to-4 mergers.

[20]Given that we are assuming a finite horizon, static CI Nash pricing every period would be the only subgame perfect equilibrium outcome under CI. However, in a infinite horizon model, joint-profit maximizing CI prices would be much higher: for example, with $N = 4$ they would be 170% higher than static Nash prices and these prices could be sustained by trigger strategies if $\beta > 0.62$.

Figure 5. : Results from a Discrete Cost Type Model with Symmetric Firms and Alternative Discount Factors.

(a) Increase in Average Price Relative to CI Nash Equilibrium

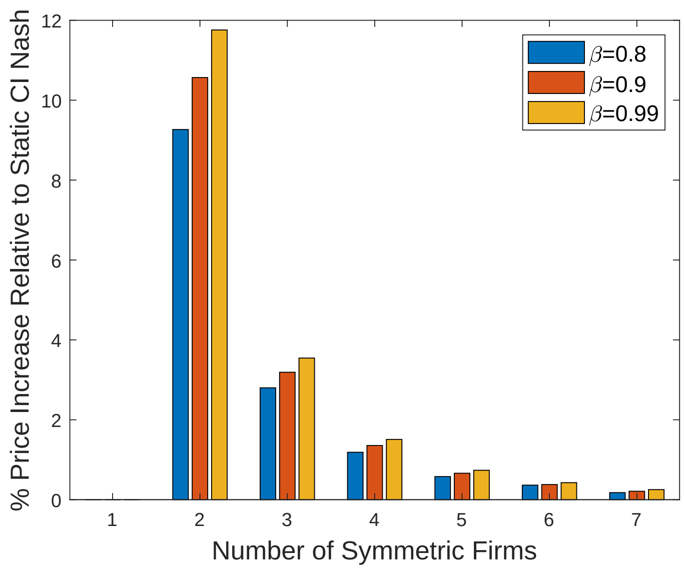

(b) Percentage Increase in Average Prices After a Merger that Eliminates a Product.

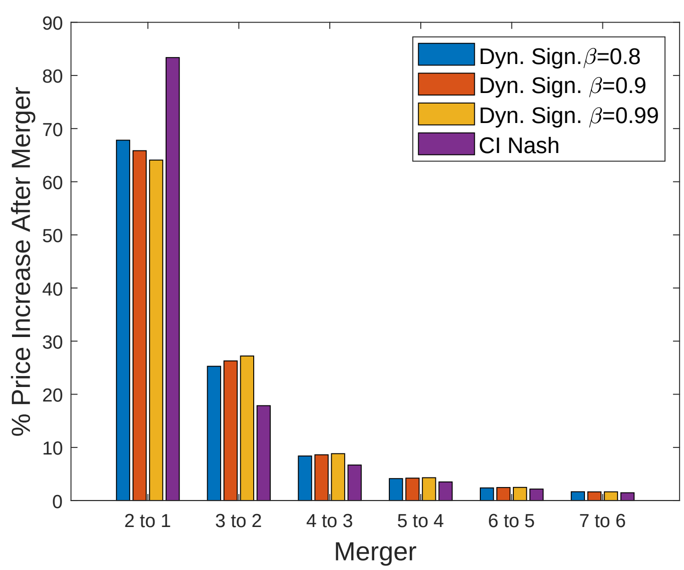

Demand Asymmetries. — We now consider how the effects of signaling vary with demand asymmetries. We assume a discount factor of 0.99, and that, as before, each firm can have a marginal cost of 8 or 8.05, with $\rho = 0.65$.[21] The nesting and price demand parameters are unchanged, but we choose the firm-specific indirect utility intercepts so that, with average marginal costs for each firm and static CI Nash prices, 97.5% of potential consumers purchase one of the products, implying limited substitution to the outside good, and the firms have shares that we specify. As before, we iterate backwards until signaling prices converge.

Figure 6(a) shows the effect of signaling on share-weighted average prices, relative to static CI Nash, with the circle areas indicating the magnitudes that are also written in the figure. The CI share of the largest firm and the split of the shares of the other firms are represented on the axes. Dynamic signaling causes a 3.4% average price increase in the symmetric 3-firm model (bottom-left circle). The percentage increases are largest when the industry is close to an effective duopoly (middle of the top row), but they can be quite large in some other cases. For example, when the CI Nash shares of total sales are {0.68, 0.24, 0.08}, average prices are 4.0% higher with dynamic signaling. In this example, it is the second largest firm whose average price increases the most (7.7%).

Figure 6(b) shows how signaling increases share-weighted average prices (across all products) after mergers in asymmetric 4-firm industries. The pre-merger model extends the example just described to an additional firm that is assumed to be symmetric with firm 1. Firm 1 and the additional firm then merge, so that the merged firm has two symmetric products after the merger. We also assume that the merged firm benefits, and is known to benefit, from the CI CMCR so that a static CI Nash model, as the agencies would use, would predict no increase in prices after the merger. The x-axis indicates the combined (CI Nash) pre-merger market share of the merging firms (so 0.5 means that each party makes 25% of sales), and the y-axis shows how the remaining shares are split between the non-merging firms.

A merger with CI CMCR efficiency in a symmetric four-firm industry (coordinates (0.5,0.5)) increases average prices by 2.3% when firms signal. Price increases are larger when one of the non-merging firms has a much larger market share (the upper rows in the figure). The intuition is that when the price set by the merged firm becomes more sensitive to its own cost and the prices set by its rivals, a dominant rival's strategy will tend to respond more, creating a greater positive feedback, than the strategies of two smaller rivals that will free-ride off each other's efforts to raise the merged firm's prices.

We also find examples where the prices of the non-merging firms rise as much or more than those of the merging firms, which would not happen in a static CI Nash model. For example, if pre-merger shares of sales are {0.325, 0.325, 0.33, 0.02}, the equilibrium average prices of the merging products rise by 7.7%, and the large rival increases its average price by 9.1%. The small rival's average price increases by 0.5%.

III. Application: MillerCoors Joint Venture

We calibrate the cost parameters of our model using data on retail prices before the 2008 Miller Coors joint venture (MCJV), and compare what our model predicts should have happened to price levels and price dynamics after the JV (an effective merger) with the changes observed in the data. The text focuses on explaining the calibration and the comparison. We use data from the IRI Academic Database (Information Resources Inc. (n.d.), Bronnenberg, Kruger and Mela (2008)). The data that we use was purchased in September 2020, and covers 2001 to 2012. Online Appendix C describes the data, and the selection choices that we make, in more detail. The appendix also

[21] If we assume a serial correlation of 0.7 or 0.75, the price effects are larger, but we are unable to find a converged fully separating equilibrium for some of the more extreme market structures we consider here.

Figure 6. : Price Effects in a Dynamic Signaling Model with Demand Asymmetries and Discrete Costs.

(a) Increase in Share-Weighted Average Prices, Relative to CI Nash, in a Three Firm Model.

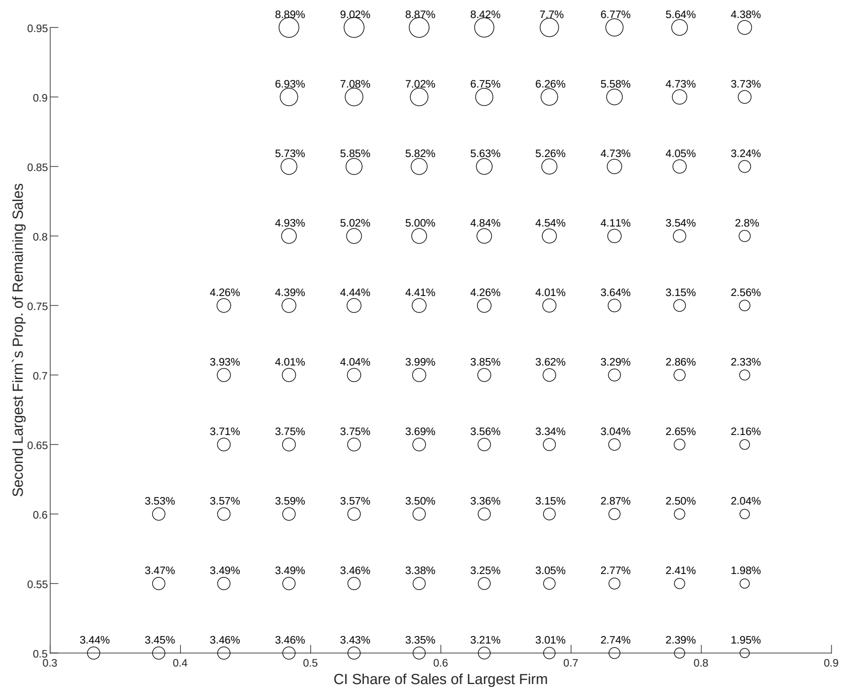

(b) Increases in Average Prices after a 4-to-3 Merger Where the Merged Firm Benefits from the CI CMCR.

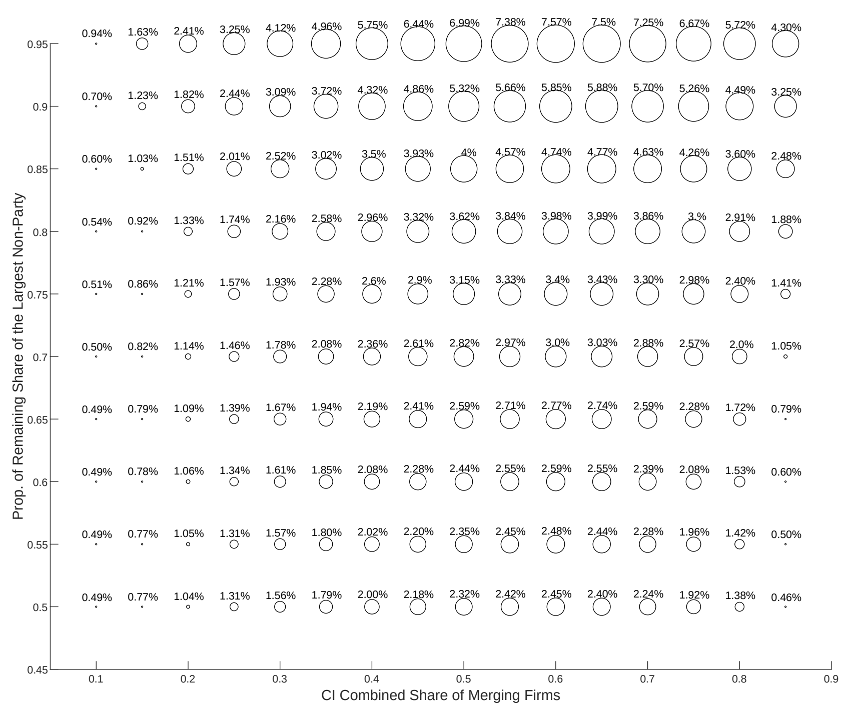

describes the types of dynamic variation in prices that identifies our parameters, and provides some additional analyses, which make use of additional data on transportation distances (Miller and Weinberg (2017*a*)), that justify our assumptions or are consistent with our interpretation of the results.

While we will show that our model predicts changes more accurately than a simple CI model with a collusive conduct parameter, we want to emphasize that we are not claiming that there are no collusive models, or models involving dynamics other than signaling, that can explain the observed changes in the data. We will return to what we believe our model potentially adds to the antitrust analysis of mergers in the conclusion.

We focus on the MCJV for two reasons. First, we want a setting where the very strong assumptions that we have to make for tractability are not too unreasonable. For example, three firms, Anheuser-Busch (AB), Miller and Coors, dominated the "subpremium" and "premium" segments of the beer industry before the JV, so that it may be appropriate to focus on only three players. After the JV, Miller Lite (ML) and Coors Light (CL) were produced in the same breweries, making our assumption that they would have the same realized costs plausible. Online Appendix C.7 shows that their post-JV prices are highly correlated, justifying our assumption that they are sold at the same price.

Second, from MW, we know that the JV was followed by rising prices of domestic brands, despite significant expected efficiencies, with AB and MC raising their prices by similar amounts (online Appendix C.3). While we know that our model can generate this outcome, we want to test whether this is the model's actual prediction when we calibrate the parameters using pre-JV data.[22]

A. *Calibration of the Dynamic Asymmetric Information Model.*

We calibrate an infinite horizon, continuous marginal cost three-firm/product version of our model to match several moments characterizing pre-JV price levels and price dynamics. As a comparison, we will calibrate the cost parameters assuming static CI Nash pricing. We treat local markets as different repetitions of the same game, rather than having different demand and cost primitives, in order to limit the computational burden.[23]

Products. — We model the pricing of three brands, which we label BL, ML and CL, although we view them as representing each brewer's broader product portfolio.[24] The products of other brewers are included in the outside good.[25]

Demand. — We assume static, time-invariant nested logit demand, with the three brands in the same nest.[26] For our baseline specification, the four demand parameters (the nesting and price parameters, and the mean utility intercepts of BL and of ML and CL, which must be the same for our counterfactual) are chosen so that, at average real pre-JV prices ($10.09 per 12-pack for BL and $9.95 for ML and CL), BL, ML and CL market shares are 28%, 14% and 14% respectively, the mean own brand price elasticity is -3 and, on average, a price increase results in 85% of the brand's lost

[22]Our approach, where we calibrate our parameters using pre-JV data, and then, assuming no change in the nature of equilibrium play, compare the model's post-JV predictions to the data, is different to the approach of MW, who use the observed post-JV increase in AB's price to estimate their collusive conduct parameter.

[23]Computationally light two-step approaches, which are often used to estimate dynamic games, require that all serially-correlated state variables, which in our setting would include beliefs, are observed by the researcher.

[24]Online Appendix C.7 shows that the pairwise correlations of prices within a portfolio are very high.

[25]An earlier version of this paper calibrated a model that included two imported brands as a non-signaling fringe in a separate nest. The model predicted that the JV would raise their prices by around 1 cent, compared with 70 cents for domestic brands.

[26]Demand-side dynamics, for example, reflecting habit formation, might also impact the effects of horizontal mergers (MacKay and Remer (2022)).

demand being diverted to the two remaining brands.[27] The implied nesting and price parameters are 0.77 and −0.10, and the BL and ML/CL mean utilities are 1.04 and 0.86 respectively. Online Appendix C.4 explains our rationale for choosing these values for the mean elasticity and diversion, but we will consider how alternative elasticity and diversion values change our results.

Marginal Costs. — The marginal costs of brand i, $c_{i,t}$, are assumed to lie on an interval $[\underline{c_i}, \underline{c_i}+c']$ and to evolve according to an AR(1) process

$$c_{i,t} = \rho c_{i,t-1} + (1-\rho)\frac{\underline{c_i} + (\underline{c_i} + c')}{2} + \eta_{i,t} \tag{5}$$

where $\eta_{i,t} \sim TRN(0, \sigma_c^2, \underline{c_i} - \rho c_{i,t-1} - (1-\rho)\frac{\underline{c_i}+\underline{c_i}+c'}{2}, \underline{c_i} + c' - \rho c_{i,t-1} - (1-\rho)\frac{\underline{c_i}+\underline{c_i}+c'}{2})$. σ_c is the standard deviation of the untruncated innovation distribution. We calibrate the five parameters $\underline{c_{BL}}$, $\underline{c_{ML/CL}}$, c', ρ and σ_c.[28]

Objective Function, Matched Statistics and Identification. — We calibrate the cost parameters using indirect inference (Smith (2008)). For a given guess, we solve the model (see online Appendices A2 and A3 for details) and simulate time-series of prices to calculate six statistics/regression coefficients. We find the parameters that provide the closest match to six similar statistics calculated from the data. The objective function to be minimized is $g(\theta)'Wg(\theta)$.

$g(\theta)$ is a vector where each element k has the form $g_k = \frac{1}{M}\sum_m \tau_{k,m}^{data} - \widehat{\tau_k(\theta)}$ where $\tau_{k,m}^{data}$ is the statistic estimated from the observed data and $\widehat{\tau_k(\theta)}$ is its simulated data equivalent. Our reported results use the identity matrix for W, although, because we match the moments almost exactly, alternative Ws give similar calibrated parameters. Minimization uses `fminsearch` in MATLAB. Standard errors are calculated treating different markets as independent observations of the same game.

The six data statistics are calculated using series of average prices from 45 regional markets from January 2001 to the announcement of the JV in October 2007. Our baseline specification uses weekly data, excluding price reductions, and the five most common pack sizes (6, 12, 18, 24 and 30-packs).[29] Market-week-brand-pack size average real prices per 12-pack equivalent are calculated using only market-week-pack sizes where we observe more than five stores carrying at least one of the flagship brands.

The first two statistics that we match are prices for BL and ML, averaged across pack sizes and weeks. The third statistic is a measure of the dispersion of BL prices, calculated as the interquartile range (IQR) of the market-specific residuals from a regression where market-week-pack size average prices of BL products are regressed on dummies for the specific set of stores observed in the market-week (interacted with pack size), to control for fixed cross-store differences in retail prices, and week-pack size dummies, to control for national promotions.

The remaining statistics are coefficients from market-brand-specific regressions where weekly brand-size prices are regressed on the lagged prices of all three brands, dummies for the specific set of stores observed in the market-week (interacted with pack size) and pack size-specific time

[27]The assumed shares overstate the share of BL relative to ML and CL, but understate the share of AB, relative to Miller and Coors, in the beer market.

[28]We have also fitted a model where ρ, σ_c and c' vary across BL and ML/CL. However, this model only improves the fit slightly, and the calibrated parameters are very similar. Moreover, when additional parameters are firm-specific it is unclear what we should assume happens to them after the JV. Our approach means that we only have to make a transparent and conventional assumption about a synergy that changes the level of marginal costs for the merged firm.

[29]Our model does not have different pack sizes, market heterogeneity, varying sets of stores or time trends, so the regressions using simulated data do not control for these factors. See online Appendix C.1 for a discussion of the sample selection.

trends. For each market, this provides estimates of nine coefficients $\rho_m^{i,j}$, where i indicates the brand of the dependent variable, and j the brand of the lagged price. We then match the average values of $\rho_m^{ML,ML}$ and $\rho_m^{CL,CL}$, $\rho_m^{BL,CL}$ and $\rho_m^{BL,ML}$, and $\rho_m^{ML,CL}$ and $\rho_m^{CL,ML}$.[30]

Assuming that there is a unique signaling equilibrium, the intuition for the identification of the cost parameters is straightforward.[31] Given the assumed demand system and observed price levels, the mark-ups implied by static best responses, which will be chosen by the lowest cost type, will identify the lower bounds on brand marginal costs.

The AR(1) price coefficients and the observed dispersion of prices will identify the range of costs and the parameters of the cost innovation process. Online Appendix C.5 illustrates the price dynamics in the pre-JV data using diagrams from the Seattle market, and regressions that are national versions of the lagged price regressions used at the market-level. It also shows how these patterns are robust to alternative ways of calculating the price series (for instance, whether temporary price reductions are included).

B. *Calibrated Parameters and Model Fit.*

Table 3 reports the calibrated parameters. Column (1) is our baseline specification. Column (7) shows the calibrated parameters for a CI Nash specification, based on the same moments. The column (1) parameters imply that $\Pr(c_{i,t+1} \geq \frac{\overline{c}+\underline{c}}{2} | c_{i,t} = \overline{c}) \approx 0.76$, which is not far from the 0.68 probability that we used in our examples, even though the support of costs is much wider. As the CI Nash model implies smaller mark-ups and less pass-through of cost changes to prices, the calibrated CI parameters imply that costs are higher on average and can vary more.

The other columns vary the assumed elasticities, diversion or price series used in estimation. Assuming more elastic demand and greater diversion to the outside good implies smaller margins in a signaling equilibrium, so the calibrated average level of costs is higher. Prices are more volative when temporary price reductions are included, so that the calibrated σ_c increases.

As discussed in online Appendix C.5, when we use monthly prices in the estimation of lagged price coefficients, we can only include market-pack size fixed effects, rather than fixed effects for the set of stores that are observed. If we include market fixed effects and use a monthly price series that excludes temporary price reductions, prices appear much more persistent. The calibrated parameters then lead the conditions required for our characterization of pricing strategies to fail in our counterfactuals. We therefore report the monthly results including price reductions.

Table 4 reports the fit of the model for the column (1), (2) and (7) specifications. The upper half of the table shows the targeted moments, and the lower half reports non-targeted moments, including the skewness of the innovations in the lagged price regressions. Skewness is a moment of a different order to the targeted moments.

The signaling models match both types of moments well, apart from tending to overpredict the IQR of price levels and underpredict the variance of price innovations, although the match of these moments is better when we only try to match moments for 12-packs. The CI model incorrectly predicts that the cross-brand ρs should be very close to zero, and it fails to match the observed skewness of price innovations.

C. *Predicted Effects of the JV.*

We resolve the model to predict the effects of the JV, and compare the predictions to changes observed in the data. In line with our simulations, we assume that, after the JV, MC owns two

[30]Our model implies that the expected $\rho_m^{ML,ML}$ should equal $\rho_m^{CL,CL}$ (etc.) so we match the average value of these coefficients.

[31]If there are multiple equilibria, minimization of the objective function may be difficult if our solution algorithm jumps between different sections of the equilibrium correspondence when the parameters are changed. In practice, we can match our chosen moments almost exactly suggesting that this is not a significant problem. We have also not found examples of multiplicity with continuous types, although we recognize that multiple equilibria may exist.

Table 3—: Calibrated Parameters for Seven Specifications.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Price Series Used for Calibration							
Data Frequency	Week	Week	Week	Week	Week	Month	Week
Pack Sizes	All	12 only	All	All	All	All	All
Temporary Price Reductions	Drop	Drop	Include	Drop	Drop	Include	Drop
Fixed Effects in Price Regressions	Stores	Stores	Stores	Stores	Stores	Market	Stores
Demand Assumptions							
Mean Price Elasticity	-3	-3	-3	-2.5	-3.5	-3	-3
Mean Diversion Between Flagship Brands	85%	85%	85%	90%	80%	85%	85%
Model	Signal	Signal	Signal	Signal	Signal	Signal	Static CI
Lower Bound Cost	\$5.248	\$5.353	\$4.856	\$4.247	\$5.973	\$4.417	\$5.405
for BL ($\underline{c_{BL}}$)	(0.195)	(0.075)	(0.328)	(0.178)	(0.230)	(0.212)	(0.084)
Lower Bound Cost for	\$6.432	\$6.588	\$6.015	\$5.800	\$6.893	\$5.572	\$6.607
ML/CL ($\underline{c_{ML/CL}}$)	(0.075)	(0.086)	(0.243)	(0.024)	(0.301)	(0.343)	(0.351)
Width Cost Interval	\$0.616	\$0.585	\$1.197	\$0.524	\$0.647	\$1.936	\$0.727
($\overline{c_i} - \underline{c_i}$)	(0.233)	(0.122)	(0.581)	(0.060)	(0.106)	(0.377)	(0.173)
Cost AR(1) Parameter	1.114	0.762	1.119	1.330	1.052	0.716	0.920
(ρ)	(0.437)	(0.213)	(0.668)	(0.145)	(0.254)	(0.408)	(0.995)
Std. Dev. Cost	\$0.268	\$0.227	\$0.573	\$0.252	\$0.272	\$0.677	\$0.278
Innovations (σ_c)	(0.098)	(0.020)	(0.319)	(0.022)	(0.087)	(0.063)	(0.164)

Note: BL = Bud Light, ML = Miller Lite and CL=Coors Light. Stores = fixed effects for the set of stores observed in the market-week (interacted with pack size) are included in the price regressions using the observed data. Standard errors in parentheses. All price series are calculated using market-week-pack sizes where at least five stores selling flagship brands are observed. Flagship diversion refers to the average proportion of lost demand that switches to the other two products when the price of one of the products increases.

Table 4—: Model Fit for Three Specifications Using Weekly Data, Average Brand Price Elasticity of -3 and Flagship Diversion of 85%.

Price Series Used for Calibration					
Data Frequency	Week	Week	Week	Week	Week
Pack Sizes	All	All	All	12-Pack	12-Pack
Temporary Price Reductions	Drop	Drop	Drop	Drop	Drop
Data/Model	Data	Signal	Static CI	Data	Signal
			Targeted Moments		
Mean p_{BL}	\$10.09	\$10.09	\$10.09	\$10.30	\$10.30
Mean p_{ML}	\$9.96	\$9.96	\$9.96	\$10.22	\$10.22
Mean $\rho^{ML,ML}, \rho^{CL,CL}$	0.389,0.395	0.391	0.391	0.468,0.450	0.458
Mean $\rho^{BL,ML}, \rho^{BL,CL}$	0.081,0.067	0.080	-0.001	0.102,0.056	0.083
Mean $\rho^{ML,CL}, \rho^{CL,ML}$	0.049,0.037	0.036	-0.002	0.065,0.026	0.040
Interquartile Range p_{BL}	\$0.141	\$0.189	\$0.198	\$0.145	\$0.186
			Non-Targeted Moments		
Mean p_{CL}	\$9.94	\$9.96	\$9.96	\$10.19	\$10.22
$\rho^{BL,BL}$	0.430	0.389	0.396	0.442	0.448
Mean $\rho^{ML,BL}, \rho^{CL,BL}$	0.057,0.0.43	0.046	0.005	0.065,0.040	0.043
Std. Dev. of Residuals					
BL	\$0.175	\$0.108	\$0.118	\$0.136	\$0.103
ML/CL	\$0.202,\$0.187	\$0.157	\$0.149	\$0.161,\$0.149	\$0.155
Interquartile Range p_{ML}, p_{CL}	\$0.168,\$0.156	\$0.272	\$0.261	\$0.169,\$0.159	\$0.277
Skewness of Residuals					
BL	-0.362	-0.339	-0.004	-0.307	-0.378
ML/CL	-0.108,-0.334	-0.325	0.010	-0.296,-0.201	-0.381

Note: BL = Bud Light, ML = Miller Lite and CL=Coors Light. Residuals are from AR(1) regressions, where the data regressions include set-of-store fixed effects. The calculation of the data statistics is explained in the text, with the model predictions based on simulating 10,000 periods of data. For the data we report separate values for the statistics for ML and CL, but, because the model assumes that ML and CL are symmetric, and so predicts identical statistics (ignoring simulation error), we match the average of these values during estimation and report a single prediction.

Table 5—: Predicted Average Prices Before and After the MC JV For Signaling Model.

	(1)	(2)	(3)	(4)	(5)	(6)
Price Series Used for Calibration						
Data Frequency	Week	Week	Week	Week	Week	Month
Pack Sizes	All	12 only	All	All	All	All
Temporary Price Reductions	Drop	Drop	Include	Drop	Drop	Include
Fixed Effects in Price Regressions	Stores	Stores	Stores	Stores	Stores	Market
Demand Assumptions						
Mean Brand Price Elasticity	-3	-3	-3	-2.5	-3.5	-3
Mean Diversion Between Flagship Brands	85%	85%	85%	90%	80%	85%
Pre-JV Mean Prices						
BL	\$10.09	\$10.30	\$9.81	\$10.09	\$10.09	\$9.75
ML/CL	\$9.96	\$10.22	\$9.68	\$9.96	\$9.96	\$9.63
Assumed ML/CL Synergy (CI CMCR)	\$1.17	\$1.20	\$1.14	\$1.50	\$0.94	\$1.13
Post-JV Mean Prices with Signaling						
BL	\$10.65 (+5.6%)	\$10.98 (+6.5%)	\$10.19 (+3.9%)	\$11.04 (+9.4%)	\$10.43 (+3.4%)	\$10.21 (+4.7%)
ML/CL	\$10.51 (+5.5%)	\$10.87 (+6.3%)	\$10.05 (+3.9%)	\$10.89 (+9.3%)	\$10.29 (+3.3%)	\$10.08 (+4.7%)

Note: BL = Bud Light, ML = Miller Lite and CL=Coors Light. Stores=fixed effects for the set of stores (interacted with pack size) observed in the market-week which are included in the price regressions using the observed data. For the data we report separate values for the statistics for ML and CL, but, because the model assumes that ML and CL are symmetric, and so predicts identical statistics (ignoring simulation error), we report a single prediction.

products, which, in each period, have the same realized marginal cost. We assume that $\underline{c_{BL}}$, ρ, c' and σ_c do not change, but we assume that $\underline{c_{ML/CL}}$ would have fallen by the synergy that would have prevented a post-JV price increase with static CI Nash pricing. The Department of Justice likely expected a synergy at least this large when it decided not to challenge the merger, and it is similar to the synergy that MW estimate.

Online Appendix C.3 shows that, after the JV, the prices of the three domestic flagship brands increased, relative to the prices of imports, by 4-6%, or 40 cents to one dollar per 12-pack. These are consistent with the estimates in MW. Table 5 shows the predicted prices for each brand before and after the JV for the six signaling model specifications. All of the predicted flagship price increases are above 3%, and the price increases for ML/CL and BL are very similar. Therefore, the predicted changes in price levels appear consistent with the data.

This broad conclusion is robust to some alternative assumptions. For example, if we assume a synergy 50 cents (42%) larger (which would lower prices in a CI Nash model), the predicted column (1) post-JV prices of ML/CL and BL are \$10.21 (a 2.5% increase) and \$10.54 (a 4.5%) respectively. On the other hand, a 50 cent smaller synergy predicts prices of \$10.79 (8.3%) and \$10.75 (6.5%) respectively. If we use a discount factor of $\beta = 0.998$, arguably more consistent with weekly data, and re-calibrate the model, the predicted post-JV prices in column (1) change by less than a couple of cents.

Figure 7. : Bud Light Equilibrium Pricing Strategies (for estimates in column (1) of Table 3).

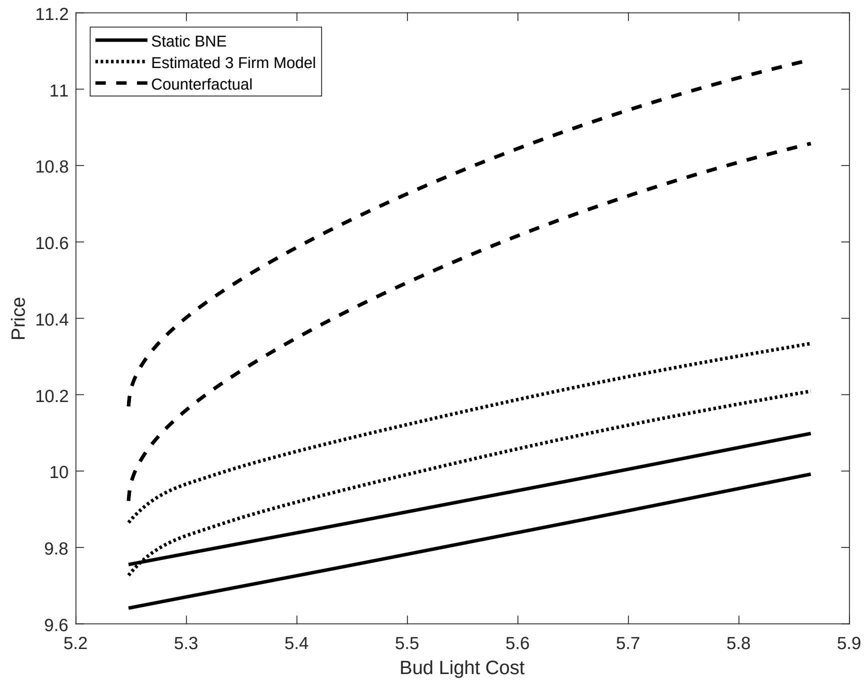

Note: the strategies shown assume that $c_{t-1}^{BL} = \underline{c^{BL}}$ and $c_{t-1}^{ML} = c_{t-1}^{CL} = \underline{c^{ML/CL}}$ (lower line) and $c_{t-1}^{BL} = \overline{c^{BL}}$ and $c_{t-1}^{ML} = c_{t-1}^{CL} = \overline{c^{ML/CL}}$ (upper line). Therefore, for each type of equilibrium, the maximum range of BL's prices spans from the lowest point on the bottom line to the highest point on the upper line.

Figure 7 compares, based on the column (1) parameters, BL's equilibrium pricing strategies for the static Bayesian Nash 3-firm model, the estimated signaling 3-firm model and the counterfactual post-JV model. Even though BL's costs are unchanged, the greater responsiveness of ML/CL pricing causes BL's prices to rise and become more volatile.

We can examine how accurately our model predicts changes in price dynamics. Table 6 compares the cross-market averages of the IQR of prices and ρ parameter statistics before and after the JV in the data, and the values predicted by the column (1) model. We also report the predicted values for the CI model when we assume that, after the JV, the firms use CI first-order conditions with a "conduct parameter" weight of 0.15 on the profits of their rivals, as such conduct gives similar predicted post-JV price increases to the signaling model.[32]

The signaling model correctly predicts the direction in which each of the reported statistics changes (i.e., post-JV prices vary more, and own-brand and cross-brand serial correlation coefficients are higher). Our ability to match these changes is encouraging given that our calibration uses no post-JV data, although we cannot rule out the possibility that other changes in the industry (e.g., changes in demand or cost shocks, or uncertainty after the 2008 financial crisis) also affect the dynamics of prices. On the other hand, the CI-conduct model predicts close to zero cross-brand ρs before and after the JV, that the dispersion of BL and ML/CL prices should move in opposite directions, and that own-brand ρs should not have changed.

[32]A value of 0.15 is smaller than MW's estimated conduct parameters, reflecting how their preferred demand model implies more substitution from domestic to imported brands than our assumed demand parameters. See online Appendix C.4.

Table 6—: Observed and Predicted Changes in Price Dynamics for Calibrated Signaling Model (Table 3, col (1)) and the Calibrated CI Model (Table 3, col (7)) with a Conduct Parameter ($\theta = 0.15$) to Predict the Same Change in Average Prices.

	Data		Calibrated Signaling Model		Calibrated CI Model with Conduct	
	Pre-JV	Post-JV (Change)	Pre-JV	Post-JV (Change)	Pre-JV	Post-JV (Change)
Interquartile Range of Prices						
BL	0.141	0.165 (+0.025)	0.189	0.356 (+0.167)	0.198	0.227 (+0.028)
ML	0.168	0.185 (+0.017)	0.271	0.355 (+0.084)	0.261	0.232 (-0.030)
CL	0.156	0.174 (+0.018)	0.271	0.355 (+0.084)	0.261	0.232 (-0.030)
AR(1) Regression Coefficients						
$\rho^{BL,BL}$	0.430	0.500 (+0.070)	0.389	0.424 (+0.035)	0.396	0.397 (+0.002)
$\rho^{ML,ML}$	0.389	0.454 (+0.065)	0.391	0.426 (+0.034)	0.391	0.393 (+0.002)
$\rho^{CL,CL}$	0.395	0.428 (+0.033)	0.391	0.426 (+0.034)	0.391	0.393 (+0.002)
$\rho^{BL,ML}$	0.081	0.100 (+0.019)	0.080	0.153 (+0.073)	-0.001	0.000 (+0.000)
$\rho^{BL,CL}$	0.067	0.096 (+0.029)	0.080	0.153 (+0.073)	-0.001	0.000 (+0.000)
$\rho^{ML,BL}$	0.057	0.098 (+0.041)	0.046	0.147 (+0.101)	0.005	0.002 (-0.003)
$\rho^{CL,BL}$	0.043	0.080 (+0.037)	0.046	0.147 (+0.101)	0.005	0.002 (-0.003)

Note: BL = Bud Light, ML = Miller Lite and CL=Coors Light. The calculation of the data statistics is explained in Section III.A, with the model predictions based on simulating 10,000 periods of data. Pre-JV averages are calculated for 45 markets, and post-JV averages are calculated for 44 markets, as one market does not have at least 5 stores observed in consecutive weeks after the JV. The CI Model simulations use the parameters from Table 3, column (7), which assumes CI Nash pricing before the JV, and assume that after the JV the firms use a conduct parameter of 0.15.

The steeper pricing functions in the counterfactual also imply an increase in the pass-through of (unobserved) marginal cost changes. For example, a 60 cent increase in BL's unobserved cost raises BL's expected pre- and post-JV prices by 40 and 85 cents respectively. While directly testing this implication is impossible without observing actual cost changes, we can identify a change in the pass-through of transportation costs which is potentially consistent with the prediction of our model.

Suppose that the unobserved portion of marginal cost reflects the trucking distance to a market, which all firms observe, multiplied by the current "per mile per unit" cost of distribution (efficiency), which is unobserved by rivals. Efficiency may vary over time depending on, for example, the capacity utilization of the brewer's truck fleet. If this is correct, the support for i's marginal cost of selling a unit of beer in market m, $[\underline{c_{i,m}}, \overline{c_{i,m}}]$, will tend to be wider for brewer-markets with longer trucking distances.

As noted in Section II, signaling will increase prices more when marginal costs have wider supports and mergers tends to make signaling effects larger. If our speculative assumptions about costs are correct, then we would expect to see larger post-JV price increases in markets with longer trucking distances. Online Appendix C.6 presents regressions that show that, in the data, post-JV price increases for domestic brands were larger in markets with longer trucking distances, including when we control for changes in market concentration and transportation synergies realized by the merger, with no significant changes for imported brands.

IV. Conclusion

Mailath (1989) showed that, in a two-period game, simultaneous signaling will increase first period equilibrium prices when firms have private information about their marginal costs. Mailath's insight has been ignored by the subsequent empirical IO and antitrust literatures, probably because it has been assumed that the extent of private information will usually be fairly small, and that limited private information would only have small effects on prices. We have shown that this second assumption is incorrect, because of feedback loops between strategies across firms, and across periods in games where firms set prices repeatedly.

As signaling incentives become stronger with fewer firms, we find that horizontal mergers can raise prices more than static CI Nash models predict even in the absence of collusion. Our model predicts observed changes in price levels, price dynamics and pass-through after the MillerCoors joint venture without requiring a change in the nature of firm behavior. Our results suggest that merger retrospectives in other industries could usefully assess whether post-merger price changes are associated with changes in dynamics and pass-through that CI Nash or simple CI conduct models cannot explain.

We have frequently been asked about the relationship between dynamic signaling and tacit collusion. One question is whether dynamic signaling should be viewed, in the language of the Horizontal Merger Guidelines, as a type of "unilateral effect" or as a type of "coordinated effect", or as a completely different sort of theory. Unilateral and coordinated effects are sometimes interpreted as reflecting those that arise in static CI Nash and CI tacitly collusive equilibria respectively (Porter (2020)). Recently, Baker and Farrell (2019) and Farrell and Baker (2021) have suggested broadening the definition of coordinated effects to include "non-purposive" theories, such as the CI Markov Perfect equilibrium models with asynchronous price-setting proposed by Maskin and Tirole (1988). We view our model as being in this class of theories, although it departs from CI while making the more conventional assumption of simultaneous price-setting.

A second question is whether there is any value in agencies or courts considering non-collusive coordinated effects theories. We believe that there are two reasons why it should be valuable. First, agencies find it very difficult to show that collusive effects are likely unless there is an established history of cartel behavior or, at least, highly suspicious parallel accommodating conduct. It might be much easier to use internal business documents to show that how firms predict their rivals' pricing is consistent with firms having incentives to signal.[33] Second, because our model predicts price effects that are usually much smaller than a tacit collusion model, it does not completely close the door on the merging parties arguing that efficiencies are likely to be large enough to counteract the upward pressure on prices.

A third question is whether there are situations in which one could more convincingly distinguish between our model and collusion as explanations for changes in price levels, price dynamics and pass-through. We have emphasized that this kind of test may be very difficult in price-setting industries where both models will tend to predict price increases, and some collusive model is likely to be able to rationalize changes in price dynamics and pass-through, but this is clearly an area where future work is important.[34] However, in a quantity-setting game where quantities are strategic substitutes, a signaling model will predict that firms will tend to overproduce, relative to

[33]In a later beer industry complaint, the Department of Justice highlighted language in AB's Conduct Plan that described its pricing as aiming to create "consistent and transparent competitive response" and to achieve the "highest level of followership" (`https://www.justice.gov/atr/case-document/file/486606/download` [accessed December 3, 2023]) as evidence of collusive intent. However, one could also view these statements as being consistent with AB wanting to make signaling as effective as possible, and it is notable that the parts of the plan that have been publicly referenced do not discuss the types of threats to punish that are an integral part of collusive strategies. In addition, as pointed out by a referee, the public record does not indicate that any of the main elements of the Conduct Plan changed after the MCJV, which is consistent with our assumption that the underlying nature of firm strategies remained the same.

[34]See Appendix D.12 of STY-WP for a more detailed analysis of how well conduct and price leadership models of collusion explain the beer data.

static CI Nash, in order to lower their rivals' future output. This is the opposite of what collusive models will predict.

REFERENCES

Ashenfelter, Orley, Daniel Hosken, and Matthew Weinberg. 2014. "Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers." *Journal of Law and Economics*, 57(S3): S67–S100.

Asker, John, Chaim Fershtman, Jihye Jeon, and Ariel Pakes. 2020. "A Computational Framework for Analyzing Dynamic Procurement Auctions: The Market Impact of Information Sharing." *RAND Journal of Economics*, 51(3): 805–839.

Baker, Jonathan B, and Joseph Farrell. 2019. "Oligopoly Coordination, Economic Analysis, and the Prophylactic Role of Horizontal Merger Enforcement." *University of Pennsylvania Law Review*, 168: 1985–2018.

Berry, Steven T, and Philip A Haile. 2014. "Identification in Differentiated Products Markets using Market Level Data." *Econometrica*, 82(5): 1749–1797.

Bresnahan, Timothy F. 1982. "The Oligopoly Solution Concept is Identified." *Economics Letters*, 10(1-2): 87–92.

Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela. 2008. "Dataset Paper: The IRI Marketing Dataset." *Marketing Science*, 27(4): 745–748.

Cardell, N Scott. 1997. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory*, 13(2): 185–213.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*, 102(2): 179–221.

Farrell, Joseph, and Jonathan B Baker. 2021. "Natural Oligopoly Responses, Repeated Games, and Coordinated Effects in Merger Analysis: A Perspective and Research Agenda." *Review of Industrial Organization*, 58(2): 1–39.

Fershtman, Chaim, and Ariel Pakes. 2012. "Dynamic Games with Asymmetric Information: A Framework for Empirical Work." *Quarterly Journal of Economics*, 127(4): 1611–1661.

Garmon, Christopher. 2017. "The Accuracy of Hospital Merger Screening Methods." *RAND Journal of Economics*, 48(4): 1068–1102.

Harrington, Joseph E. 2021. "There May Be No Pass Through of a Merger-Related Cost Efficiency." *Economics Letters*, 208: 110050.

Information Resources Inc. n.d.. "IRI Academic Data Set 2001-2012." Information Resources Inc., Chicago, Illinois. Data purchased September 2020.

Kaya, Ayça. 2009. "Repeated Signaling Games." *Games and Economic Behavior*, 66(2): 841–854.

Lau, Lawrence J. 1982. "On Identifying the Degree of Competitiveness from Industry Price and Output Data." *Economics Letters*, 10(1-2): 93–99.

MacKay, Alexander, and Marc Remer. 2022. "Consumer Inertia and Market Power." Harvard University.

Mailath, George J. 1988. "An Abstract Two-Period Game with Simultaneous Signaling: Existence of Separating Equilibria." *Journal of Economic Theory*, 46(2): 373–394.

Mailath, George J. 1989. "Simultaneous Signaling in an Oligopoly Model." *Quarterly Journal of Economics*, 104(2): 417–427.

Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand curves, and Edgeworth Cycles." *Econometrica*, 571–599.

Mester, Loretta J. 1992. "Perpetual Signalling with Imperfectly Correlated Costs." *RAND Journal of Economics*, 23(4): 548–563.

Milgrom, P., and J. Roberts. 1982. "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis." *Econometrica*, 50(2): 443–459.

Miller, Nathan H, and Matthew C Weinberg. 2017*a*. "Supplement to "Understanding the Price Effects of the MillerCoors Joint Venture"." Econometric Society [distributor], https://doi.org/10.3982/ECTA13333. Accessed March 9, 2019.

Miller, Nathan H, and Matthew C Weinberg. 2017*b*. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica*, 85(6): 1763–1791.

Nevo, Aviv. 1998. "Identification of the Oligopoly Solution Concept in a Differentiated-Products Industry." *Economics Letters*, 59(3): 391–395.

Olley, G Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263–1297.

Peters, Craig. 2009. "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry." *Journal of Law and Economics*, 49(2): 627–649.

Porter, Robert H. 2020. "Mergers and Coordinated Effects." *International Journal of Industrial Organization*, 73: 102583.

Roddie, Charles. 2012. "Signaling and Reputation in Repeated Games, I: Finite Games." University of Cambridge.

Shapiro, Carl. 1986. "Exchange of Cost Information in Oligopoly." *Review of Economic Studies*, 53(3): 433–446.

Smith, Anthony. 2008. "Indirect Inference." *The New Palgrave Dictionary of Economics, 2nd Edition.*

Sweeting, Andrew, James W Roberts, and Chris Gedge. 2020. "A Model of Dynamic Limit Pricing with an Application to the Airline Industry." *Journal of Political Economy*, 128(3): 1148–1193.

Sweeting, Andrew, Mario Leccese, and Xuezhen Tao. 2023. "Should We Expect Merger Synergies to Be Passed Through to Consumers?" University of Maryland.

Sweeting, Andrew, Xuezhen Tao, and Xinlu Yao. 2022. "Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers." National Bureau of Economic Research No. 28589.

Sweeting, Andrew, Xuezhen Tao, and Xinlu Yao. 2023. "Replication Code for "Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers"." Inter-university Consortium for Political and Social Research [distributor], http://doi.org/10.3886/E192405V1.

Toxvaerd, Flavio. 2008. "Strategic Merger Waves: A Theory of Musical Chairs." *Journal of Economic Theory*, 140(1): 1–26.

Toxvaerd, Flavio. 2017. "Dynamic Limit Pricing." *RAND Journal of Economics*, 48(1): 281–306.

Vives, Xavier. 2011. "Strategic Supply Function Competition with Private Information." *Econometrica*, 79(6): 1919–1966.